

SAM (REX/Ops) Operations Meeting 2008.06.03

Adam (Typist), Art, SteveW, SteveS, Angela, Randolph, Robert, Eric, Anil
Meeting time: 55 minutes

For reports submitted to the meeting...

The report text is in this font

Comments made during the meeting are in this font inline with the report

Releases:

Randolph has availability status changes in CVS. Robert has a request in to Steve to fix samadmin cache cleanup commands to access the history table. Also note that if you delete a disk, you will lose its history information. Perhaps need a data_disk history table too. Maybe do with a trigger instead of modifying the DB server. Perhaps wait for all of these changes before release. ETA is "a week or so".

Also, thread safety code is not as safe as it should be (for the client side) for running python threads.

DØ:

--- May 28: More than 1 report complained that enstore was very bad that day.

Twice heard complaints that only one mover was up, but when we looked later, things seemed ok. Check with enstore-admin. Robert is pre-caching the ichep data.

--- Is <https://fermilab.onjira.com/login.jsp> the place to login ? I couldn't login myself.

This is because we have not attached the LDAP server to JIRA (doing so requires a security analysis from FNAL computer security - that should happen this week). There is a local JIRA password set up - if you don't know what it is, go to the main JIRA page (fermilab.onjira.com) and click on "Forgot Password". Enter your e-mail address and it should send you the password. If not, let me (Adam) know.

--- There doesn't seem to be minutes for the last few meetings ? I was hoping to possibly look for answers to my questions.

I'll post what I have backlogged.

Unresolved issues:

#3900: MC jobs held : I can't login to JIRA to see.

Looking into this.

#3899: http://www-d0en.fnal.gov/enstore/enstore_saag.html seemed to show many yellow buttons for a long time (from yesterday to today ?).

#3894: cannot copy file from sam : CSskim-EMinclusive-20060108-125957-1806424.fi x_p17.07.01 doesn't exist though SAM still has a pnfs location. Should we delete that location or ??

This file was recycled - but still has a pnfs entry. Robert is cleaning these up.

#3890: never ending jobs : few OSG MC jobs that report all local jobs have finished but the grid job still reports running (after a long time). Can this be fixed ?

Still looking at this one.

CDF:

=====
Monday 7PM Chicago (Tuesday May 27 9 AM Seoul)

(Presumably) Known Problems from DHAAG page:

1. No of cells Inactive: 4: r-fcdfdata320-{1,4}
2. No of cells OFFLINE : 7: r-fcdfdata154-{1,4}, data157-{1,3}
3. Bad Movers: 9940B22, 26

Tuesday 3PM Chicago

Email from Tim Messer

"The LTO3 library managers for CDFEN and STKEN are paused because of a problem with a robot arm in the SL8500 unit at GCC.

A service call has been placed with STK.
I will send updates as warranted."

Tuesday 8:45 PM Chicago (10:45 AM Seoul)
- As of 10:45(8:45PM Chicago),
CDF-LTO3 seems back online.
(CDF-9940B has been online since 6:15PM Chicago)
- A Bad Tape IAE661 in CDF-LTO3 showed up.
an email sent to enstore-admin ccing cdfdh_oper
- 2 Bad Movers in CDF-LTO3 as well.
will keep on an eye on it for a few more hours

Tuesday 9:50 PM Chicago
Total number of files to restore are high (3800) now
but appear to decrease gradually as enstore is back online.

Wednesday 2:42 AM Chicago
A bad mover in CDF-LTO3 LTO3_23 for >4 hours.
An email sent to enstore-admin.
Also a bad tape IAB446 in CDF-LTO3, reported to enstore-admin

Wednesday 1 PM Chicago
Email from Tim Meser
"LTO3 libraries on CDFEN and STKEN are paused due to an LSM failure
on the GCC SL8500 robot. A service call has been placed with Sun/STK. "

Wednesday 1 PM Chicago
An email sent to enstore-admin
"I just wanted to let you know that there are 7 bad movers
and one bad tape in CDF-LTO3."

Wednesday 2:45 PM Chicago
Email from Rick Snider
"I just spoke to Nelly and have concluded that we need to proceed
with the downtime to fix database index corruption (on Thursday)"

Wednesday 4:16 PM Chicago
Email from Tim Messer
SL8500 robot arm has been replaced, and the LTO3 libraries on both
CDFEN and STKEN are back in service.

Thursday 2:05 AM Chicago
All of CDF-LTO3 movers are marked as "bad"
and all transfers in LTO3 are "pending"..
Sent an email to enstore-admin around midnight.
CDF ACE emailed cdfdh_oper saying that
"ENCP process is red in ProcMon since about five hours ago" around 1AM.
I just talked with CDF ACE and heard that this problem
will not affect CDF data taking even for the next 6 hours
when enstore people come to office around 8AM.
i.e. I am not sure whether to call SAM oncall shifter now(2AM)
and decided to wait until morning..

Thursday 7:04 AM Chicago
Email from George Szmuksta
There is another problem with LSM 0,1 in the SL8500 at GCC a service
call will be placed.

Friday 8:45 Chicago
A bad mover LTO3_05. email to enstore-admin

During the daily meeting with AB, Angela noticed a problem
and discussed with Krzysztof.

Email from AB

"There was a jump in the plot of the difference between restore requests and queued jobs on the **CDF** dCache Wednesday morning. I suspected the usual NOACCESS tape, but upon further investigation that seems to not be the case. I have tracked down a few entries that exist on the restores page but not on the queues page, and they all seem to think they are running on r-fcdfdata136-4. However, that pool has no movers, active or queued, on it. I'm not sure what to do at this point. Are these requests essentially "lost" in this state? Do we need to track down which job(s) they come from and kill the jobs?

Reply from KG

"If it is not caused by a tape, then this may be an internal dcache discrepancy. This was one of the past reasons why this plot was created.

One way to resolve it may be to restart the pool after moving it to TestPools and having it drained (this includes the staging requests), but it may not help. Then one can restart the PoolManager. This usually causes the stage requests to be retried though so be careful with it as it may double the number of requests to enstore.

In any case I would contact dcache-admin as they may want to look at the situation before anything is changed to see if it is not one of the bugs one should fix."

First problem fixed itself Friday afternoon (when there was a corresponding jump in the bpool's jobs). Happened again Sunday morning with 20 request difference. Disappeared Monday at noon. Both times asked Alex to look into it - haven't heard anything back. Two open JIRA tickets on this. Perhaps this is some internal dCache corruption.

Friday 2:45 PM Chicago (4:45 AM Seoul)

Email from Tim Messer

I will be taking a few cdfen movers down for a short time this afternoon to upgrade them to SL4 and SDE. These will be taken down one at a time to minimize disruption:

cdfenmvr101a (LTO3_01)
cdfenmvr102a (LTO3_02)
cdfenmvr104a (LTO3_04)
cdfenmvr105a (LTO3_05)
cdfenmvr110a (LTO3_10)

Reponse from AB

We still have a lot of requests pending for using those drives (built up from the recent Enstore down times). Could this work be delayed until Monday?

Sunday 2 AM Chicago (Monday 4 PM Seoul)

a new bad mover, 9940B16
and a bad tape, IA4030.
email to enstore-admin

End of Hyunwoo Kim shift 5/27 - 6/2

Yesterday - from ntupler processing (Ray Culbertson) - using fcp to transfer files to the upload server. Was checking fcp once per minute. Was keeping a count on amount of data transfered, would increase about 1 Gig/minute but saw a big drop of 100 Gig within 5 minutes - and files were not transfered. Under investigation (CDFDH-101).

=====

MINOS:

General

Global BlueArc failures Sunday 1 June , 05:30 to 07:30, and 10:03

Monitoring picked this up,
<http://www-numi.fnal.gov/computing/dh/bluwatch/>

KCA upgrade Wed 28 May delayed farm processing,
the robot cert was not preregistered.

Enstore -

DCache -

Jun 10 FNDCA public DCache upgrade to 1.8 - is this scheduled ?
Jun 19 FNDCA public DCache downtime for SDE all day.
It does not seem reasonable to have two all-day outages so close.
Please combine them.

SAM -

SAM-IT/3582 - provided login and testing instructions to developers.
Status ?

Oracle -

Oracle quarterly patches - pending - probably June 10 (or June 19 ?)

OPEN ISSUES

IT 1146: cannot test existence of sam tape location
IT 2071: quadratic project startup time growth

PLAN -

DB:

D0 :

Alert(s):
- Message=Tablespace [OPERATIONAL_IDX_01] is [80 percent] full

Action: Added 6GB

CDF :

- Oracle's Quarterly patch on cdfopr2 database. - 29th May - Done
- Recreated the indices on cached_files table.

Alert(s):

MINOS:

- Scheduling Oracle's Quarterly patch on minosprd June 10th

dCache/enstore:

enstore/dcache report

Scheduled Downtimes:

o d0en - Tue, Jun 3, 7:30 am - 11:30 am

ADIC needs to replace the arm again since it can't reach some of the IO boxes (used to work). Will need a July downtime.

- ADIC AML/2 maintenance
- enstore s/w minor update

o fndca - Thu, Jun 10, 8 am - 4 pm

- public dcache s/w upgrade to release 1.8

o stken - Thu, Jun 19, 8 am - 4 pm

- public enstore s/w upgrade to SDE

Operations:

o SL8500s

- Repeated problems starting Tuesday with one of the bots in SL8500 #1, affecting stken and cdfen, appears to have been resolved by Thursday. Resolution was delayed because remedy failed to page anyone.

o d0en

- All LTO2 drives were idle for a period Wednesday afternoon, resolved by restarting inquisitor service. We have no explanation.
- 18 new LTO4 movers are installed and have been checked out with their drives. They should be available tomorrow after the downtime.

Qizhong needs to be alerted when these drives are ready.

o cdfen

o cdfdca

- 4 read pools on fcdldata154 are offline.
- 3 write pools on fcdldata157 are offline.

o stken

- 2 LTO4 drives had been down for several weeks; problem resolved by swapping fibers to drives.
- stkensrv3 was rebooted to fix a problem with the raid array

o fndca